GENERATIONS / VANCOUVER 12-16 AUGUST

SIGGRAPH2018

12th AUGUST

WORKSHOP

COMPUTER GRAPHICS FOR AUTONOMOUS DRIVING APPLICATIONS

EXTENDED PROGRAM

GENERATIONS / VANCOUVER 12-16 AUGUST
SIGGRAPH2018

12th AUGUST

WORKSHOP | PROGRAM

COMPUTER GRAPHICS FOR AUTONOMOUS DRIVING APPLICATIONS

8:45 h	Workshop Introduction. Opening
9:00 h	**Jose M Alvarez** (NVidia Corp.) *Scaling up Deep Learning for Autonomous driving*
9:35 h	**Jose De Oliveira** (Unity Tech.) **Miguel Ferreira** (CVedia) *Environment and Sensor Simulation for Autonomous Vehicle Training*
10:10 h	**Simon Box** (Aurora) *Graphics for visualization, testing and training self-driving vehicle development*
10:30 h	Coffee Break
11:00 h	**Dinesh Manocha** (UMD) *Simulation and Navigation for Autonomous Vehicles*
11:35 h	**Ashu Rege and Yongjoon Lee** (Zoox) *Advanced 3D Simulation for Autonomous Vehicles*
11:55 h	**Panel Discussion I**
12:30 h	Lunch (Catering)
13:30 h	**Ming C. Lin** (UMD) *From Learning-based Traffic Reconstruction to Autonomous Driving*
14:05 h	**Kevin McNamara** (Parallel Domain) *Building the Virtual Road to Autonomy*
14:25 h	**Philipp Slusallek** (DFKI, Saarland University) *Understanding the World via Digital Reality: Research at the Crossroads of Visual Computing, Artificial Intelligence and High Performance Computing*
15:00 h	Coffee Break
15:30 h	**Germán Ros** (Intel Corp.) *CARLA: Towards Democratizing Simulation in Autonomous Driving*
16:05 h	**Gavriel State** (NVidia Corp.) *Training Deep Networks with Synthetic Data: Bridging the Reality Gap*
16:25 h	**Panel Discussion II**
17:00 h	Closing Remarks.

12th AUGUST

WORKSHOP | PROGRAM

COMPUTER GRAPHICS FOR AUTONOMOUS DRIVING APPLICATIONS

Scaling up Deep Learning for Autonomous driving 9:00

Jose M. Alvarez (NVidia Corp.)

Deep learning has rapidly moved from research to be a key component in providing industrial impact in areas such as autonomous driving. From initial semantic segmentation to more recent advanced systems, algorithms continue consuming data and computational resources. These data-driven methods are eager in consuming data and computational resources. Data being acquired and the need of annotations keep growing exponentially and open new challenges to improve accuracy and to achieve the desired safety level. In this talk, I will explore some of these challenges along with our proposed solutions in terms of computational efficiency or the efficient use of synthetic data for training deep networks.

PROFILE

Jose M. Alvarez is a Senior Deep Learning Engineer at NVIDIA working on scaling-up deep learning for autonomous driving. Previously, he was a senior researcher at Toyota Research Institute and at Data61/CSIRO (formerly NICTA) working on deep learning for large scale dynamic scene understanding. Prior to that, he worked as a postdoctoral researcher at New York University under the supervision of Prof. Yann LeCun. He graduated with his Ph.D. from Autonomous University of Barcelona (UAB) in October 2010, with focus on robust road detection under real-world driving conditions. Dr. Alvarez did research stays at the University of Amsterdam (in 2008 and 2009) and the Electronics Research Group at Volkswagen (in 2010) and Boston College. Since 2014, he has served as an associate editor for IEEE Transactions on Intelligent Transportation Systems.

Environment and Sensor Simulation for Autonomous Vehicle Training 9:35

Jose De Oliveira (Unity Tech.) + Miguel Ferreira (CVedia)

One of the biggest pain points for Autonomous Vehicle (AV) development is the generation of content to train and validate machine learning models. Real-world data collection is an expensive process that most time falls short of the multiple scenario requirements required for AV training. Synthetic environments and real-time 3D rendering provide an effective alternative. In this talk, we will show how Unity and CVEDIA have been leveraging the capabilities of the platform to create multi-sensory environment simulations for autonomous applications enabling us to enrich real-world datasets with synthetic data, identify and reproduce edge cases, and increase model performance and confidence.

PROFILE

Jose has 20 + years of industry, working at tech giants such as IBM, Microsoft and Uber in areas that range from real-time communications to enterprise security systems. In 2006 he focused his career on Machine Learning, working on content filtering solutions for Family Safety at Microsoft, where he headed the delivery of the first SmartScreen anti-phishing solution for Internet Explorer 7. He later drove the development of paid search relevance models for mobile devices at Bing Ads and worked on applying Machine Learning to geospatial problems at Bing Maps, continuing that work after joining Uber in 2015. In 2017 he joined the Machine Learning Team at Unity, leading the autonomous vehicles engineering project, part of Unity's Industrial initiatives. He's based out of Bellevue, WA.

12th AUGUST

WORKSHOP | PROGRAM

COMPUTER GRAPHICS FOR AUTONOMOUS DRIVING APPLICATIONS

PROFILE

MIGUEL FERREIRA
After helping brands, like Ferrero, De Agostini, MindChamps, shape the mobile entertainment space, he is now Senior Software Engineer at CVEDIA, where he pushes the boundaries of real-time rendering, developing sensor models for cutting-edge deep learning applications. SynCity is a hyper-realistic simulator specifically designed for deep learning algorithm development and training. Constructing complex 3D land, aerial and marine environments and generating ground truth data for sensor devices like LiDAR, radar, thermal, near and far IR, and cameras, SynCity unleashes the limitations of the physical world. When Miguel is not simulating the real world, he is traveling it in search of the perfect picture.

Graphics for visualization, testing and training self-driving vehicle

10:10

Simon Box (Aurora)

This talk provides a high level overview of the uses of graphics in self-driving car development from the perspective of Aurora Innovation. Firstly the various uses of graphics for visualization are addressed. In particular how visualization is employed in the development of different parts of the autonomy stack. Secondly, graphics for synthetic sensor input data are discussed in the context of closing the control-loop on self-driving vehicle simulation. Finally, applications of graphics for generating synthetic training data for machine learning models are considered. From bootstrapping for rare cases through correcting biases in collected data to examples of fully synthetic training. Each of these cases are illustrated with examples from the literature and from work at Aurora.

PROFILE

Simon Box is the Simulation Architect at Aurora Innovation, where the sim team is working to build a simulation framework that can virtually prototype all parts of the Aurora self-driving software stack. Simon's previous work in simulation includes his PhD at the University for Cambridge, UK where he was simulating the trajectories of particles in electrostatic fields. In the Machine Learning and Perception group at Microsoft Research, where he build a rocket flight simulator and on the Autopilot team at Tesla Motors, where he led the simulation efforts.

12th AUGUST

WORKSHOP | PROGRAM

COMPUTER GRAPHICS FOR AUTONOMOUS DRIVING APPLICATIONS

Simulation and Navigation for Autonomous Vehicles 11:00

Dinesh Manocha (University of Maryland)

In this talk, we give an overview of our recent work on simulation and navigation technologies for autonomous vehicles. This includes our work on autonomous simulator, AutonoVi-Sim, which is built on top of UnReal Engine and utilizes many modules available in game engines. AutonoVi-Sim is a collection of high-level extensible modules which allows for the rapid development and testing of vehicle configurations, and facilitates construction of complex road networks. Autonovi-Sim supports multiple vehicles with unique steering or acceleration limits taking into account vehicle dynamics constraints.

In addition, AutonoVi-Sim supports navigation for non-vehicle traffic participants such as cyclists and pedestrians AutonoVi-Sim also facilitates data analysis, allowing for capturing video from the vehicle's perspective, exporting sensor data such as relative positions of other traffic participants, camera data for a specific sensor, and detection and classification results. We highlight its performance in traffic and driving scenarios. We also describe new algorithms for planning and navigation in dense road conditions. Our approach takes into account the shapes and dynamics of road entities like cars, pedestrians, bicycles, trucks and use them to design local navigation methods. We also infer driver behavior's from the vehicle trajectories and use them to design safe navigation strategies. We also demonstrate their application on datasets available via Baidu Apolloscape. We also highlight many open issues in this area and highlight the potential of using concepts from computer graphics, virtual reality, computer vision and robotics for autonomous driving simulators.

PROFILE

Dinesh Manocha is the Paul Chrisman Iribe Chair in Computer Science & Electrical and Computer Engineering at the University of Maryland College Park. He is also the Phi Delta Theta/Matthew Mason Distinguished Professor Emeritus of Computer Science at the University of North Carolina – Chapel Hill. He has won many awards, including Alfred P. Sloan Research Fellow, the NSF Career Award, the ONR Young Investigator Award, and the Hettleman Prize for scholarly achievement. His research interests include multi-agent simulation, virtual environments, physically-based modeling, and robotics. His group has developed a number of packages for multi-agent simulation, crowd simulation, and physics-based simulation that have been used by hundreds of thousands of users and licensed to more than 60 commercial vendors. He has published more than 500 papers and supervised more than 35 PhD dissertations. He is an inventor of 9 patents, several of which have been licensed to industry. His work has been covered by the New York Times, NPR, Boston Globe, Washington Post, ZDNet, as well as DARPA Legacy Press Release. He is a Fellow of AAAI, AAAS, ACM, and IEEE and also received the Distinguished Alumni Award from IIT Delhi. He was a co-founder of Impulsonic, a developer of physics-based audio simulation technologies, which was acquired by Valve Inc.

12th AUGUST

WORKSHOP | PROGRAM

COMPUTER GRAPHICS FOR AUTONOMOUS DRIVING APPLICATIONS

Advanced 3D Simulation for Autonomous Vehicles 11:35

Ashu Rege and Yongjoon Lee (Zoox)

How do we build the Matrix for autonomous vehicles to prove themselves and learn from within? We will discuss the challenges involved with using simulation for safety validation and machine learning, as well as Zoox's approach to various aspects including determinism, variety, realism, and scalability of synthetic worlds.

PROFILE

Ashu Rege is the Vice President of Software at Zoox, responsible for Zoox's entire software platform including machine learning, motion planning, perception, localization, mapping, and simulation. Ashu joined Zoox from NVIDIA, where he was VP of Computer Vision & Robotics responsible for NVIDIA's autonomous vehicle and drone technology projects. Previously, he held other senior roles at NVIDIA including VP of the Content & Technology group developing core graphics, physics simulation and GPU computing technologies, and associated software. Prior to NVIDIA, he co-founded and worked at various startups related to computer graphics, laser scanning, Internet and network technologies. Ashu holds a Ph.D in Computer Science from U.C. Berkeley.

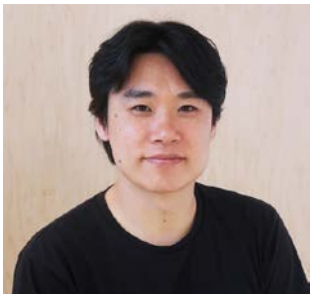

Yongjoon Lee is Engineering Manager of Simulation at Zoox, responsible for the simulation platform to validate and improve the safety and quality of autonomous driving software. Yongjoon joined Zoox from Bungie, where he worked as engineering lead for AI, animation, core action system, cinematic system, and mission scripting system teams. Prior to Bungie, he published six technical papers at SIGGRAPH on realistic motion synthesis using reinforcement learning. He holds a Ph.D in Computer Science & Engineering from the University of Washington.

12th AUGUST

WORKSHOP | PROGRAM

COMPUTER GRAPHICS FOR AUTONOMOUS DRIVING APPLICATIONS

From Learning-based Traffic Reconstruction to Autonomous Driving

13:30

Ming C. Lin (University of Maryland)

Rapid urbanization and increasing traffic have led to digitalization of modern cities and automation of transportation means. As new technologies like VR systems and self-driving cars emerge, there is an increasing demand to incorporate realistic traffic flows into virtualized cities. In this talk, we first present a novel method for learning-based traffic animation and visualization using GPS data. This approach reconstruct city-scale traffic using statistical learning on GPS data and metamodel-based simulation optimization for dynamic data completion in areas of insufficient data coverage. Next, we present a unified collision avoidance algorithm for the navigation of arbitrary agents, from pedestrians to various types of robots, including vehicles.

This approach significantly extends our statistically-based WarpDriver algorithm specialized for disc-like agents (e.g. crowds) to a wide array of robots in a unifying framework using different nonlinear motion extrapolations of motion to support agent dynamics, with additional shapes and soft constraints, to simulate vehicle traffic. Finally, we introduce a learning-based, multi-level control policy for autonomous vehicles by analyzing simulated accident data and using our collision avoidance algorithm, data annotation, and parameterized traffic & vehicle simulation. We conclude by suggesting possible future directions.

PROFILE

Ming C. Lin is currently the Elizabeth Stevinson Iribe Chair of Computer Science at the University of Maryland College Park and John R. & Louise S. Parker Distinguished Professor Emerita of Computer Science at the University of North Carolina (UNC), Chapel Hill. She is also an honorary Chair Professor (Yangtze Scholar) at Tsinghua University in China. She obtained her B.S., M.S., and Ph.D. in Electrical Engineering and Computer Science from the University of California, Berkeley. She received several honors and awards, including the NSF Young Faculty Career Award in 1995, Honda Research Initiation Award in 1997, UNC/IBM Junior Faculty Development Award in 1999, UNC Hettleman Award for Scholarly Achievements in 2003, Beverly W. Long Distinguished Professorship 2007-2010, Carolina Women's Center Faculty Scholar in 2008, UNC WOWS Scholar 2009-2011, IEEE VGTC Virtual Reality Technical Achievement Award in 2010, and many best paper awards at international conferences. She is a Fellow of ACM, IEEE, and Eurographics.

12th AUGUST

WORKSHOP | PROGRAM

COMPUTER GRAPHICS FOR AUTONOMOUS DRIVING APPLICATIONS

Building the Virtual Road to Autonomy 14:05

Kevin McNamara (Parallel Domain)

Real-world training and testing of autonomous vehicles leads to significant safety risks and logistical challenges. Over the course of millions of miles, even a tiny risk of collision becomes inevitably large. As a result, the technology behind video games and movies is powering state of the art simulations aimed at improving the safety and reliability of self-driving cars. But if we need to drive millions (and potentially billions) of miles in a simulator, how do we build those miles? Constructing high fidelity 3D worlds and scenarios is prohibitively slow, expensive, and difficult. Parallel Domain's automatic world and scenario generation technology is revolutionizing the way autonomous vehicle companies scale up their simulation efforts and is paving the way for these vehicles to get it right in simulation before they hit the road.

PROFILE

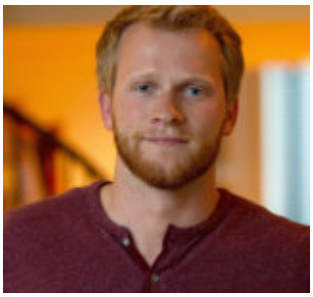

Kevin is the founder and CEO of Parallel Domain, a fast growing startup which has automated the generation of high fidelity virtual worlds and scenarios for simulation. He brings deep computer graphics experience having built and led a team within Apple's Special Projects Group focused on autonomous systems simulation, architected procedural content systems for Microsoft Game Studios, and contributed to academy award winning films at Pixar Animation Studios. Kevin holds a degree in computer science from Harvard University and resides in Palo Alto, CA.

Understanding the World via Digital Reality: Research at the Crossroads of Visual Computing, Artificial Intelligence, and High-Performance Computing 14:25

Philipp Slusallek (DFKI, Saarland University)

Artificial Intelligence (AI) systems that deal with reality need to reliably make accurate decisions even in highly complex and critical situations -- in particular when human lives are at stake. However, suitable real training data is often hard to come by, especially for critical situations that hardly ever happen.

Digital Reality is an approach to simultaneously learn models of the real world and then use them for the training of AI systems by synthetically generating the needed data via simulations. This approach also allows us to systematically validate AI systems by automatically generating test cases from our models that specifically address critical aspects. Validation with real data also allows us to continuously identify limitations in the models and adapt them to the dynamic changes in the real world. For HPG a specific focus is on open questions in predictive rendering of sensor data, especially for non-optical sensors like radar.

12th AUGUST

WORKSHOP | PROGRAM

COMPUTER GRAPHICS FOR AUTONOMOUS DRIVING APPLICATIONS

PROFILE

Philipp Slusallek is Scientific Director at the German Research Center for Artificial Intelligence (DFKI), where he heads the research area on Agents and Simulated Reality. At Saarland University he has been a professor for Computer Graphics since 1999, a principle investigator at the German Excellence-Cluster on "Multimodal Computing and Interaction" since 2007, and Director for Research at the Intel Visual Computing Institute since 2009. Before coming to Saarland University, he was a Visiting Assistant Professor at Stanford University. He originally studied physics in Frankfurt and Tübingen (Diploma/M.Sc.) and got his PhD in Computer Science from Erlangen University. He is associate editor of Computer Graphics Forum, a fellow of Eurographics, a member of acatech (German National Academy of Science and Engineering), and a member of the European High-Level Expert Group on Artificial Intelligence. His research covers a wide range of topics including artificial intelligence, simulated/digital reality, real-time realistic graphics, high-performance computing, motion modeling & synthesis, novel programming models, computational sciences, 3D-Internet technology, and others.

CARLA: Towards Democratizing Simulation in Autonomous Driving

15:30

German Ros (Intel Corp.)

Almost one year has passed since we released CARLA, an open-source simulator aiming to foster research and development of autonomous driving systems. During this year, we have experienced an increasing amount of support and feedback from the community, leading to a platform that is robust and widely used for tackling driving related problems, such as learning driving policies using reinforcement learning RL and imitation learning IL. In this talk we introduce CARLA's main features and describe what are the challenges ahead that must be solved in order to accommodate needs from both academia and industry. We also present how CARLA has been used by the research community in the context of policy learning and virtual-to-real policy transfer.

PROFILE

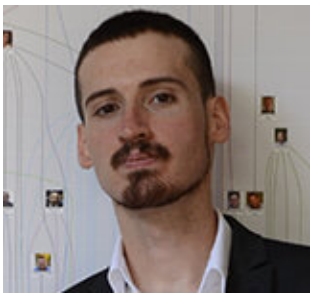

German Ros is a Research Scientist at Intel Intelligent Systems Lab (Santa Clara, California), working on topics at the intersection of Machine Learning, Simulation, Virtual worlds, Transfer Learning and Intelligent Autonomous agents. He leads the CARLA organization as part of the Open Source Vision Foundation. Before joining Intel Labs, German served as a Research Scientist at Toyota Research Institute(TRI), where he conducted research in the area of Simulation for Autonomous Driving, Scene Understanding and Domain Adaptation, in the context of Autonomous Driving.He also helped industrial partners, such as Toshiba, Yandex, Drive.ai, and Volkswagen to leverage simulation and virtual words to empower their machine learning efforts and served at the Computer Vision Center (CVC) as a technical lead for the simulation team German Ros obtained his PhD in Computer Science at Autonomous University of Barcelona & the Computer Vision Center.

12th AUGUST

WORKSHOP | PROGRAM

COMPUTER GRAPHICS FOR AUTONOMOUS DRIVING APPLICATIONS

Training Deep Networks with Synthetic Data: Bridging the Reality Gap

16:05

Gavriel State (NVidia Corp.)

In this talk we will cover the use of synthetic data in training deep neural networks for computer vision tasks. We will explain why this is a critical research area with application ranging from robotics to autonomous vehicles, and we will discuss some important techniques for generating synthetic data using domain randomization and 3D graphics engines.

PROFILE

Gavriel State is a Senior Director, System Software at NVIDIA, based in Toronto, where he leads efforts involving applications of AI technology to gaming and vice versa, in addition to work in remastering games for NVIDIA's SHIELD TV platform. Previously, Gav founded TransGaming Inc, and spent 15 years focused on games and rendering technologies.

12th AUGUST

WORKSHOP | PROGRAM

COMPUTER GRAPHICS FOR AUTONOMOUS DRIVING APPLICATIONS

WORKSHOP ORGANISERS

Jose A. Iglesias-Guitián (Computer Vision Center - UAB)

Jose A. Iglesias-Guitian is a Research Scientist in Visual Computing and Computer Graphics. He has recently joined the Computer Vision Center (CVC) at the Universitat Autonoma de Barcelona (UAB). He is currently collaborating with international companies on projects on the intersection of graphics with computer vision and autonomous driving. Before he was an Associate Researcher at Disney Research (The Walt Disney Company), based in Edinburgh (UK).

Jose has also spent research periods at the University of Zaragoza (Spain) and for more than five years he was at the Visual Computing Group of CRS4 (Italy) where he received a Ph.D. Degree in Electronic and Computer Engineering (2011). Jose has obtained predoc and senior EU Marie Curie fellowships. His research is focused on real-time graphics, volumetric and translucent material rendering, noise reduction and xR interactive applications. More information about his research work is available here (www.j4lley.com).

Antonio M. López (Computer Vision Center - UAB)

Antonio M Lopez (www.cvc.uab.es/~antonio/) is Associate Professor at the Computer Science department of the Universitat Autonoma of Barcelona (UAB). He is also a founding member of the Computer Vision Center (CVC) of the UAB, where he created and is the Principal Investigator of the group on ADAS and Autonomous Driving since 2002. Antonio is founding member and co-organizer of consolidated international workshops such as the Computer Vision in Vehicle Technology (CVVT) and the Transferring and Adapting Source Knowledge in Computer Vision (TASK-CV). Antonio has also collaborated in numerous projects with international companies, especially from the automotive sector. His work in the last 10 years has focused on the use of Computer Graphics to train and test On-board Computer Vision systems. He is the responsible at CVC/UAB of well-known simulation systems such as CARLA and SYNTHIA.

12th AUGUST

WORKSHOP | PROGRAM

COMPUTER GRAPHICS FOR AUTONOMOUS DRIVING APPLICATIONS

Remember to tweet using the hashtag
#SIGGRAPH2018 and mention
@SIGGRAPH and **@CVC_UAB** !

More information:
https://s2018.siggraph.org

Computer Vision Center